

CHAPTER EIGHT

Speciation, Phylogeny, and the Evolution of Metacommunity Biodiversity

In the 142 years since publication of *The Origin of Species* (Darwin 1859), there has been remarkable progress in evolutionary biology and paleobiology in understanding natural selection, the genetical basis of evolutionary change, and the history of evolution. As a result of this progress, there remains today no credible scientific alternative to the conclusion that species did indeed evolve through a process of organic evolution. Ironically, however, much less progress has been made since Darwin on the *origin* question itself—understanding the origin of species and characterizing the differences that, at a fundamental level, define and identify species. A recent U.S. National Academy of Sciences colloquium (NAS 1997) in honor of Theodosius Dobzhansky was nominally dedicated to new approaches to the origin of species. But, although there were many papers on the origin and maintenance of genetic differences among populations and species, few of the papers had the temerity to broach the subject of *origins* directly.

There is good reason for this hesitation because the “species concept” has varied and sometimes conflicting meanings in different scientific communities (Otte and Endler 1989; McKinney and Drake 1993). However, putting speciation into the theory of island biogeography forces a head-on confrontation with the question of what is actually being counted by parameter ν , the probability of a “specia-

CHAPTER EIGHT

tion event" per birth. Here I will argue that the mere process of asking this question—in the formal context of the unified neutral theory—potentially leads to a profound change in perspective on species, speciation, and the meaning of biodiversity itself. In this chapter I first examine the implications of the neutral theory for phylogeny and phylogenetic reconstruction. I then take a deeper look at the speciation process itself, and the implications of the theory for the nature of species and of biodiversity in general. In this chapter I also examine a second mode of speciation, the "random fission" mode, and show how this alters the equilibrium distribution of relative species abundance in the metacommunity.

Recent advances in molecular biology have opened up new possibilities for understanding phylogenetic relationships far better than ever before. These advances have stimulated renewed theoretical interest in phylogenetic reconstruction (Hey 1991, 1992, Hey et al. 1993, Scotland et al. 1994, Nee et al. 1994, Harvey and Nee 1994). The predictions of the unified theory are qualitatively different in a number of important regards from contemporary phylogenetic theory, however. In current theory, the phylogenetic unit is the *lineage*, and lineages are assigned probabilities of speciating and going extinct—probabilities that have been treated as either constants or stochastic variables by different authors. In the unified neutral theory, however, the phylogenetic unit is the *individual*. Lineages per se do not have assigned probabilities of birth and death. Instead, the stochastic rates of lineage origination and extinction are predicted by the fundamental biodiversity number, θ , and follow inevitably from the distribution of metacommunity relative species abundances. Metacommunity size J_M and the speciation rate ν dictate not only the distribution of lineage lifespans but also the potential of lineages to speciate over their evolutionary lifetimes. Thus, one significant difference from current phylogenetic theory is the fundamental importance attached by the unified theory to the relative abundance of

METACOMMUNITY BIODIVERSITY

the species that are the lineages of the metacommunity over evolutionary time.

In 1979 marine ecologist Jeff Levinton published an important but insufficiently noticed paper in *Science* that pre-saged the unified neutral theory of biodiversity and biogeography. Levinton presented a simple verbal model to explain how a diversity equilibrium might arise in the fossil record in the absence of niche assembly. The importance of Levinton (1979) is best understood retrospectively, in my opinion, in the context of the debate that ensued after publication of the "punctuated equilibrium" theory (Eldridge and Gould 1972). Eldridge and Gould argued that the evidence from the fossil record is less consistent with a process of gradual clade diversification in many taxonomic groups, and more consistent with a process of episodic diversification, such that between bouts of rapid cladogenesis, there were relatively long periods of diversity equilibrium. As noted in chapter 1, however, a diversity equilibrium does not necessarily imply a lack of turnover of species (Patzkowsky and Holland 1997). In any event, this theory stimulated the study of neutral models of randomly branching phylogenetic clades. The idea was to find out whether such models could reproduce the patterns of lineage diversification observed in the fossil record.

Raup et al. (1973) constructed a simple stochastic birth-death process as a neutral model for the branching of phylogenetic trees. They assumed that the per lineage probability of a "birth," or the division of a lineage into two daughter lineages, λ , was the same from one lineage to another, and likewise, that the probability of the death of a lineage, μ , was the same from lineage to lineage. Although total clade extinction is possible when the birth rate exceeds the death rate ($\lambda > \mu$) since it is a stochastic process, the more usual outcome in this case is geometric growth in the number of lineages within clades (Nee et al. 1994). This raises a serious problem with Raup et al.'s (1973) pure birth-death model

CHAPTER EIGHT

and other similar models, namely that they do not predict a nonzero diversity steady state (Gould et al. 1977, Stanley 1979, Raup 1985). A diversity equilibrium can be achieved in these models only if additional rules are imposed on how speciation and extinction depend on the number of lineages, a kind of diversity density dependence (Raup et al. 1973, Rosenzweig 1975; see also below).

The general failure of these constant birth-death neutral models to reproduce patterns of steady-state diversity between punctuational events, and other discrepancies between observed and predicted phylogenies, led many to reject neutral models prematurely. Before Levinton (1979), the general conclusion was that the apparent diversity equilibria seen in the fossil record must imply the existence of adaptive community diversity equilibria due to the filling of all available niche space. This means, so the argument goes, that the metacommunity must be niche assembled, and the number of niches sets a limit on total steady-state species richness in the metacommunity. This hypothesis suffers from the same problem that afflicts contemporary, niche-assembly theories of relative species abundance, viz., a free parameter (the number of available niches) that cannot be derived from the theory's first principles.

The significance of Levinton's (1979) paper was to demonstrate that a simple neutral model was fully capable of explaining diversity steady states in the fossil record without invoking any niche assembly rules whatsoever. Imagine a widespread species or taxon that is the stem ancestor to a large clade of descendant species. Assume that this ancestral species and all its contemporary descendant species compete in a zero-sum game for limiting resources. Levinton did not use the term "zero-sum game," but he clearly established an equivalent condition that all occupiable space (or limiting resources) be saturated with organisms. This ancestral species or taxon would have been common with an extensive global range and thus would have been very resistant

to extinction due to its very large metapopulation size. Now, start an engine of speciation, by whatever mechanism. New species arise from the ancestral species as well as from species descended from the original stem ancestor. As more and more species are added to the metacommunity, given the zero-sum rule, the average population size of extant species in the clade must decline. As average population size falls, the extinction rate will inevitably rise. A point will be reached when there are so many rare species that the rate of extinction in the clade will increase until it is equal to the rate of origination of new species. At this point the living biodiversity of the clade will increase no further, and a steady-state species richness will be achieved. Levinton's argument was based on average population sizes, not on the steady-state relative abundance distribution of species predicted by the fundamental biodiversity number θ —which of course was unknown at the time—but nonetheless Levinton's conclusion is qualitatively completely correct.

Levinton's model was not the only one to produce an equilibrium diversity. Rosenzweig (1975, 1995) proposed a model of the evolution of continental diversity in which the number of species approaches a steady-state asymptote. Rosenzweig (1995) argued that increased diversity per se will cause increased extinction rates per species. If the unified neutral theory is correct, however, increased extinction rates are simply an incidental by-product of the reduced mean abundances of species when there are more species in the metacommunity. In Rosenzweig's model, unlike Levinton's, the fundamental connection of the diversity equilibrium with limiting resources was not made. Under zero-sum dynamics, average population size must decline with increased species richness. In Levinton's model, as in the unified theory, species per se do not saturate the landscape and limiting resources, *individuals* (collectively) do (unless $\theta = \infty$). Of course, individuals make up species, but

from the perspective of the zero-sum game, what matters is the number of individuals, not the number of species.

Rosenzweig also argued that increased diversity will reduce the speciation rate per species. Once again, according to the unified neutral theory, lower speciation rates are not caused by higher diversity but are incidentally correlated negatively with diversity. In the theory, the number of species originating per unit time is controlled by the size of the metacommunity J_M and the per capita speciation rate ν . As the number of species increases for fixed J_M and ν , the speciation rate per species must necessarily decline. Thus, a major difference between the unified neutral theory and previous theories is that *the number of new species arising per unit time is a function of the total number of individuals in the metacommunity, not the number of pre-existing species*. A steady-state diversity is achieved, not because the total metacommunity speciation rate falls. It is achieved because the extinction rate increases until it equals the speciation rate, which is an inevitable outcome of a falling mean population size of metacommunity species under the zero-sum rule.

This perspective puts a different interpretation on the punctuated equilibrium theory itself and, particularly, on the rapid diversification in many taxa that followed mass extinctions in the fossil record. If the unified theory is correct, then a major effect of a mass extinction, beyond killing of a large fraction of the Earth's biota, is to massively desaturate the Earth's resources for a brief period, geologically speaking. Given this brief desaturation, the unified theory predicts that each mass extinction event will be followed by an episode of very rapid population expansion among the relative small number of species that survive the event. During this short nonequilibrium, post-extinction population flush, there will be a very large net excess of births over deaths—many more births per unit time than during periods of diversity equilibrium. In the unified theory, the opportunity for speciation is a function of the expressed birth rate. The

METACOMMUNITY BIODIVERSITY

theory asserts that it is primarily the temporary but huge surplus of births over deaths that generates the sudden burst of new species following mass extinctions. Such a burst of speciation is predicted *even if there is no change* in the probability of speciation per birth.

Thus, there is no a priori necessity to invoke niche assembly arguments to explain punctuated equilibrium patterns in the fossil record, although adaptive radiation may also occur as a correlated effect of the ecological release accompanying a desaturation of resources. Not all periods of rapid diversification of particular taxa occurred in the aftermath of mass extinctions. Even in these cases, however, the theory suggests that we should seek evidence of regional or global desaturation in the resources limiting the taxon in question precisely during periods of rapid diversification in the taxon. This desaturation of resources can result from the rapid decline or extinction of one or more particularly common metacommunity species in the taxon, or from a sudden increase in limiting resources available to the taxon. The theory also predicts that there should be a "dosage effect," a correlation between the degree of resource desaturation and the amount of subsequent diversification that occurs. Desaturation can also arise from a sudden, massive increase in metacommunity size J_M , which will in turn cause inflation of the fundamental biodiversity number θ . Therefore, diversification will occur with increased J_M even in the complete absence of any change in the probability of speciation per birth ν .

Before discussing the unified theory's predictions about phylogeny, it is useful to outline the current neutral theory of phylogenetic reconstruction. As mentioned, the pioneering work on neutral models in phylogeny was done in the 1970s, particularly by Raup et al. (1973) and Gould et al. (1977). These models have been further developed theoretically by Nee et al. (1992, 1994, 1995), and their predictions have been compared with several molecular phylogenies,

CHAPTER EIGHT

particularly of birds (Sibley and Ahlquist 1990), by Harvey and Nee (1994), and Harvey et al. (1994b, 1996). In the case of phylogenies derived from molecular data, the only observable species are extant species, so the full phylogeny as described in the model of Raup et al. (1973), including all its extinct lineages, is unknown. How do one's conclusions differ if one has a "censored" phylogeny, consisting of only its living members? Nee et al. (1994) distinguished four cases based in part on the degree to which the complete or incomplete phylogeny is knowable; figure 8.1 shows a cartoon of the two cases that need concern us. Case 1 represents the full phylogeny from the pure continuous-time birth-death process. In this process, lineages give rise to new lineages at a per lineage rate of λ and go extinct at a per lineage rate of μ , and the clade survives for an arbitrary length of time—the process studied by Raup et al. (1973). Case 2 (case 4 of Nee et al. 1994) is the same phylogeny as case 1, but it retains only those lineages that lead to species alive at the present time T , and from which all extinct lineages




FIG. 8.1. Two cases of phylogenies. Case 1: Phylogeny resulting from the pure birth-death process, with all lineages, living and extinct shown. Case 2: The same clade as in case 1 but with all extinct lineages removed, leaving only lineages alive at the present time T .

are pruned out. Case 2 trees are what Nee et al. (1994) call a *reconstructed phylogeny*, because it is the only phylogeny that can be inferred from molecular data on living taxa. Nee et al. (1994) make the assumption that all reconstructed phylogenies are monophyletic and that reticulate evolution (e.g., allopolyploidy) is unimportant (but see Arnold 1997).

Following Nee et al. (1994), let $\Pr\{i, t\}$ be the probability that a clade has i lineages at time t , and let $\Pr\{t, T\}$ be the probability that a single lineage alive at time t has at least one surviving lineage at present (time T). Now consider the pure birth-death process (case 1). $\Pr\{t, T\}$ is equal to:

$$\Pr\{t, T\} = \frac{\lambda - \mu}{\lambda - \mu \exp\{-(\lambda - \mu)(T - t)\}}.$$

Thus, $\Pr\{0, T\}$ is the probability that a clade starting with a single lineage at time 0 is not extinct at time T . The probability that a clade starting with a single lineage is extinct at time t is therefore

$$\Pr\{0, t\} = 1 - P(0, t),$$

and the probability that the clade has i lineages at time t is given by

$$\Pr\{i, t\} = \Pr\{0, t\}(1 - u_t)u_t^{i-1}, \quad i > 0,$$

where

$$u_t = \frac{\lambda\{1 - \exp[-(\lambda - \mu)t]\}}{\lambda - \mu \exp[-(\lambda - \mu)t]}$$

is the probability of a new lineage arising in the time interval 0 to t . $\Pr\{i, t\}$ has a simple geometric distribution (Kendall 1948).

The pure birth-death process cannot be observed directly because of past extinctions within the clade, so we need to derive the equivalent expressions for the "censored" phylogeny consisting of the lineages surviving to the present. Nee et al. (1994) show that the probability of i observed

lineages surviving to time T (case 2) is also geometrically distributed with a different geometric parameter: u'_t of the pure birth-death process (case 1) weighted by the ratio of the probabilities of a single lineage surviving to the present (time T) and to intermediate time t :

$$\Pr\{i, t, T\} = (1 - u'_t)(u'_t)^{i-1}, \quad i > 0,$$

$$\text{where } u'_t = u_t \frac{P(0, T)}{P(0, t)}.$$

Note that there is no zero term in the expression for $\Pr\{i, t, T\}$. This is because at least one lineage will survive to the present. Otherwise we generally would not know of the clade's existence; but even if it were known, we would not have the opportunity to collect molecular data from it.

The predicted shape of the frequency distribution of number of the descendant lineages from the geometric model of Nee et al. (1994) is shown in figure 8.2. In all cases, the most frequent category represents lineages with only one living descendant (themselves), and the frequency of lineages with a higher number of descendants falls off exponentially. The longer the time period sampled, the larger the number of possible descendant lineages, and the frequency distribution becomes flatter; but the highest frequency category remains the singleton category.

Some phylogenies have distributions like figure 8.2, but many do not. Some lineages are far too "bushy" for the geometric distribution. Harvey and Nee (1995) and Harvey et al. (1995, 1996) conclude that these nonfitting phylogenies are inconsistent with neutral cladogenesis. According to the unified theory, however, this conclusion is premature. The geometric distribution is *not* the distribution predicted by the unified neutral theory—except in cases of lineage-poor clades or small time samples. If the time window is too short or if the taxonomic range is too narrow, then most of the clade will not be sampled, and the distribution will

METACOMMUNITY BIODIVERSITY




FIG. 8.2. The geometric distribution of the number of descendant lineages expected under pure birth-death process of Nee et al. (1994). Arbitrary time units birth rate $\lambda = 0.0001$ and death rate $\mu = 0.00005$. *Black-bar histogram*: Phylogeny sampled over a period of 5000 time units. *Gray-bar histogram*: Phylogeny sampled over 20,000 time units. Note flattening of the distribution as the sampling interval increases, as the frequency of larger numbers of descendant lineages increases. In the geometric model, the category of one descendant lineage is always the most frequent.

appear to be geometric-like. However, the unified theory predicts that as longer and longer time windows are considered, an interior mode will appear in the distribution of number of descendant lineages, and the most abundant category will no longer be singletons.

This effect is superficially analogous to the “veil line” sampling phenomenon discovered by Preston (1948) in the distribution of relative species abundance (chapter 2). In spite of superficial similarity, however, the interior mode in the distribution of the number of descendant lineages has a different explanation than the interior mode of the relative species abundance distribution (chapter 5). The unified theory’s explanation follows from the prediction that individual lineages must have very different probabilities of

CHAPTER EIGHT

speciating and going extinct because they differ enormously in their abundances over evolutionary time. Species that are abundant and widespread throughout the metacommunity have very long expected lifespans measured in geological time, far longer than the lifespans of rare and local species. Besides having much shorter lifespans, rare lineages have fewer opportunities to speciate also because they produce absolutely fewer offspring per unit time, even though their per capita birth rate is the same. Therefore, common species will be the ancestors of more species and of more present-day biodiversity than rare species.

The actual distribution of the number of descendant lineages will also be affected by the frequency distribution of abundances of ancestral lineages. From metacommunity relative species abundance under point mutation speciation, we know that there are fewer common lineages than lineages of intermediate abundance, and fewer intermediate-abundance lineages than rare lineages. Multiplying longevity (which is a function of lineage abundance) by the number of lineages of a given abundance results in most speciation events occurring in ancestral lineages having intermediate abundances. This is the explanation for the interior mode of the distribution of the number of descendant lineages. An interior mode is also found under random fission speciation, but the metacommunity relative abundance distribution in this case is zero-sum multinomial, so that rare species are less frequent and even less likely to be the ancestors of many descendant lineages (see below).

Distributions of the expected number of descendant lineages are shown in figure 8.3 as a function of the abundances of the ancestral lineages and of the fundamental biodiversity number, θ . The distributions are negatively skewed, with most of the daughter lineages (percentage of speciation events) coming from abundant ancestral lineages. When θ is small, there is high dominance and low diversity in the metacommunity, and the distribution of speciation events

METACOMMUNITY BIODIVERSITY

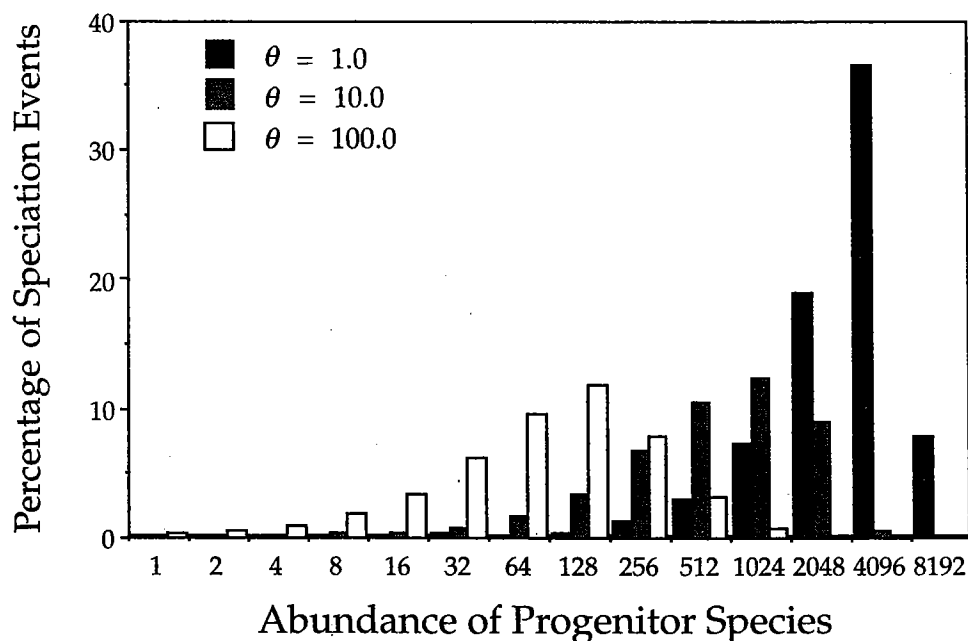


FIG. 8.3. Distribution of the number of daughter lineages at steady-state diversity equilibrium in the metacommunity, as a function of the fundamental biodiversity number θ . Note the negative skew of the distributions and the concentration of daughter lineages in high-abundance ancestors.

is right shifted. As θ increases, there are more metacommunity species of intermediate abundance, so the distributions become broader and lower. However, the distributions remain negatively skewed, such that most of the steady-state diversity always originates from the more abundant ancestral species, not from the rarer species.

As an empirical example, the phylogeny of birds represented by the distribution of species per monophyletic family clearly has an interior mode (fig. 8.4). These data were tabulated by Sibley and Ahlquist (1990) and include 8501 of the estimated 9700 extant species of birds, and more species than the 1700 species included in Sibley and Ahlquist's partial molecular phylogeny of the birds.

Herein, in my opinion, lies a fundamental conceptual weakness in the neutral models of Raup et al. (1973) and Nee et al. (1994, 1995)—and all other current models of phylogeny of which I am aware—namely, their failure to




FIG. 8.4. Distribution of the number of species per family of birds, for 8501 of the approximately 9700 species of birds in the world's avifauna. Data compiled by Sibley and Ahlquist (1990). The line is the best-fit lognormal.

take the relative abundance of lineages into account. Nee et al.'s results were derived under the assumption that λ and μ are probabilities that can be measured and defined on a per lineage basis. In contrast, the argument of Levinton (1979) and the unified theory is that the instantaneous per lineage probabilities of speciation and extinction are dictated by the current relative abundances of lineages. The omission of relative abundance is surprising because, at least in demography theory, population size has long been recognized as *the* critical parameter controlling the expected time to extinction (Levins 1970, Richter-Dyn and Goel 1972, Lande 1988, 1993).

One possible explanation for this omission is that paleobiology utilizes species presence-absence data more often than data on relative species abundance. Thus, the importance of relative abundance is perhaps less widely appreciated in

paleobiology than in population and community ecology. It may also be that the quality of relative abundance data in the fossil record is more difficult to assess. However, I suspect that the data have simply not been collected in many cases. Indeed, data on relative lineage abundance through time in the fossil record might actually be a better predictor of relative lineage longevity than the short-term data on modern relative species abundance available to ecologists.

In the unified neutral theory, we can study analytically how factoring in relative lineage abundance affects the theory of phylogeny. First we need to compute the expected lifespans of species undergoing zero-sum drift in the metacommunity as a function of their initial abundance. This problem is closely related to the "absorbing case" problem that we studied in chapter 4, namely the expected time that it takes a species undergoing zero-sum ecological drift to go locally extinct or achieve monodominance. In the present metacommunity problem, the absorbing state of monodominance no longer exists because of the low but finite probability of a speciation event. The only remaining absorbing state is global extinction of the species from the entire metacommunity. Therefore the matrix of transition probabilities is identical to the absorbing case studied in chapter 4, with the exception that

$$\Pr\{J - 1|J\} = \nu$$

and

$$\Pr\{J|J\} = 1 - \nu.$$

That is, the probability per birth that a monodominant species in the metacommunity will be dethroned from monodominance is the probability of a speciation event ν .

With this change, we can now calculate exactly the expected lifespan of arbitrary species i at initial abundance N_i , where lifespan is measured in total number of deaths in the community until extinction of the i th species, using the

same method detailed in chapter 4. The mean lifespan Ω of species i starting with N_i individuals is

$$\Omega(N_i) = \sum_{k=1}^{N_i-1} k \frac{\binom{J_M}{k}}{\binom{J_M-2}{k-1}} + N_i \left[\sum_{k=N_i}^{J_M-1} \frac{\binom{J_M}{k}}{\binom{J_M-2}{k-1}} + \frac{1}{\nu} \right].$$

The first summation is the number of times that the i th species, starting at abundance N_i , will pass through abundances 1 through $N_i - 1$ before extinction, and the second summation is the number of times that the i th species will pass through abundances N_i through $J_M - 1$. The final term, N_i/ν , is the time i th species will spend being monodominant before extinction.

Note that the dominant term in $\Omega(N_i)$ is N_i/ν , the time the i th species spends in the monodominant state. This is because leaving the state of monodominance must await a speciation event, which can be a very long wait indeed. Because this is so very much longer than the residency times at other abundances, and because real species have a vanishingly small probability of ever achieving monodominance in the entire metacommunity, at least for large metacommunities, I will henceforth ignore this term. Figure 8.5 shows how the initial abundance of a species affects its residency time at each abundance $1 \leq N_i \leq J_M - 1$ before the species goes extinct. I have illustrated the behavior of $\Omega(N_i)$ for a very small metacommunity ($J_M = 16$), but the results are qualitatively the same for all larger metacommunity sizes. Species starting their existence at small population sizes have short residency times at all abundances. Species starting at high abundance have progressively longer residency times at higher abundances.

One of the most interesting special cases is the extinction time of newly arisen species. Under the point mutation mode of speciation, calculated as probability ν per birth, the

METACOMMUNITY BIODIVERSITY

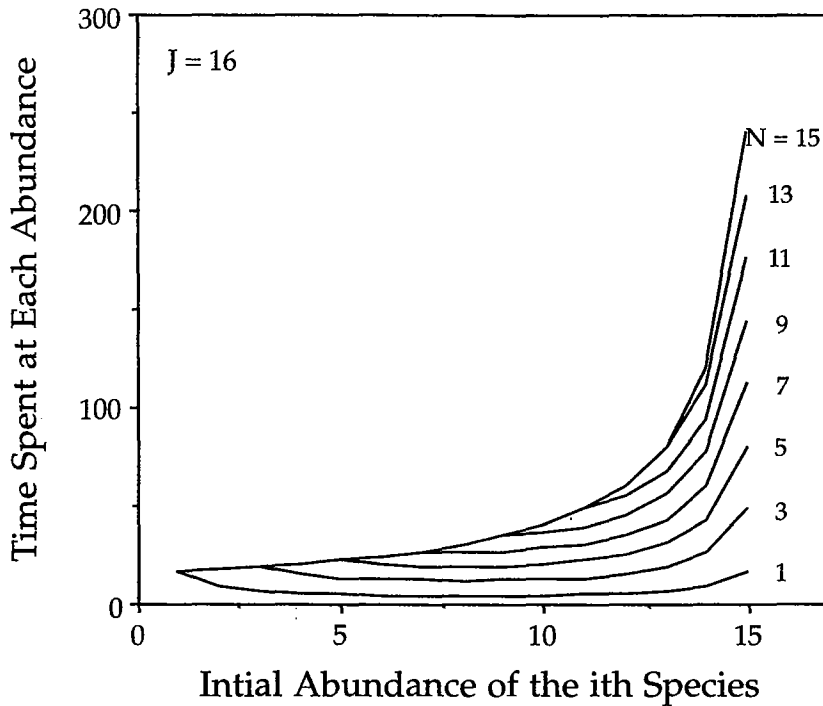


FIG. 8.5. Time spent at each abundance (residency time), short of complete monodominance, as a function of the initial abundance of the i th species, in a community of size 16.

expected time to extinction (total number of deaths in the metacommunity) of a newly originated species is (ignoring monodominance)

$$\Omega(1) \cong \sum_{k=1}^{J_M-1} \frac{\binom{J_M}{k}}{\binom{J_M-2}{k-1}}.$$

As noted in chapter 4 regarding time to fixation, calling $\Omega(N_i)$ an “extinction time” is a bit misleading because it actually represents the number of deaths in the metacommunity before a species starting with abundance N_i goes extinct. $\Omega(N_i)$ counts deaths in *all* species, not just deaths in species i . The value of $\Omega(N_i)$ increases with metacommunity size mainly because more deaths of all species occur per unit time in a large metacommunity than in a small one.

CHAPTER EIGHT

Converting the number of deaths until extinction to an absolute time scale is actually straightforward. All we need to do is normalize for metacommunity size. For example, on average twice as many deaths occur per unit time in a metacommunity that is twice as large. Therefore, to put extinction times in these two communities on the same absolute timescale, simply divide the number of deaths by two in the community that is twice as large. After normalization, however, one will still find that a species starting from a given abundance will live longer in a larger metacommunity. This is because a species starting with N_i individuals in a larger metacommunity has potentially more states of abundance that it can pass through before extinction occurs. Figure 8.6 shows how total time to extinction of the i th species depends on initial population size and the size of the metacommunity. In this figure I have scaled all extinction times on the same timescale, relative to the number of deaths in the smallest community ($J_M = 16$). Thus, the total number of




FIG. 8.6. Scaled extinction times for metacommunities of different sizes, as a function of the initial abundance of the i th species.

deaths before extinction of the i th species for the community of size 32 is divided by 2; the total for the community of size 64 is divided by 4, and so on.

We are now in a position to calculate the distribution of lifespans expected at equilibrium between speciation and extinction in the metacommunity. Recall from chapter 5 that the expected metacommunity relative abundance distribution is given by

$$E\{r_i|J_M\} = \sum_{k=1}^C r_i(k) \cdot \Pr\{S, r_1, r_2, \dots, r_S, 0, 0, \dots, 0\}_k,$$

where C is the total number of configurations, $r_i(k)$ is the abundance of the i th ranked species in the k th configuration, and $\Pr\{S, r_1, r_2, \dots, r_S, 0, 0, \dots, 0\}_k$ is the probability of the k th configuration. The metacommunity dominance-diversity curve is the set of ordered expectations, $E\{r_i\}$, $i = 1, 2, \dots$, ordered such that species of the lowest rank are the commonest. Therefore, the equilibrium distribution of lifespans is the distribution of

$$\Omega(E\{r_i\}), i = 1, 2, \dots$$

In general, the $E\{r_i\}$ are nonintegral expectations. To calculate this distribution, one can use the gamma function to calculate nonintegral factorials, at least for species whose expected abundances are greater than unity. I prefer an alternative, however. In chapter 9, I present the recipe for stochastically simulating the metacommunity relative abundance distribution. Each simulation results in only integer abundances of each species. For each simulation one can compute $\Omega(N_r)$ exactly, where N_r is the observed abundance of the r th ranked species in the given simulation. Then, to obtain arbitrarily accurate estimates of the distribution of lifespans, one averages across the ensemble of simulation results. Accuracy increases as a function of sample size, the number of simulation runs.

CHAPTER EIGHT

Interestingly, the fundamental biodiversity number θ does not uniquely determine the distribution of species lifespans, i.e., the distribution of Ω , even though θ does uniquely determine the steady-state metacommunity distribution of relative species abundance. However, the two parameters that compose θ , metacommunity size J_M and speciation rate ν , both separately enter the function for Ω .

I now compute the equilibrium distribution of lifespans of all metacommunity species in a metacommunity of size $J_M = 100,000$. Figure 8.7 displays what might be termed *lifespan-diversity curves* in analogy to dominance-diversity curves. The x -axis is species rank in abundance, as in a standard dominance-diversity curve. However, the y -axis is the common logarithm of the mean lifespan for species of a




FIG. 8.7. Lifespan-diversity curves measured in total deaths in a metacommunity J_M of size 100,000, for various values of the fundamental biodiversity number θ . Means of 1000 simulations. The breaks in the curves occur at the rank above which all communities cease having the same number of species. As $J_M \rightarrow \infty$, the curves approach straight lines as in the logseries for infinite metacommunity size.

METACOMMUNITY BIODIVERSITY

given rank in abundance in the metacommunity. The curves are for three values of the fundamental biodiversity number θ : 1, 10, and 100. As $J_M \rightarrow \infty$, the curves approach a logseries-like straight line on the semilog plot of log extinction time versus species rank, and the downward inflection point disappears. This downward bend in the curve is found in finite metacommunity sizes and arises because not all communities of a given size have the same number of species. The larger the metacommunity, the higher the rank at which the inflection point occurs, for a given θ value. Note that increasing metacommunity diversity actually decreases the lifespans of the most common species, although it increases the lifespans of species of intermediate abundance (fig. 8.7).

The most important conceptual point made by Figure 8.7 is that there are enormous differences in longevity among the metacommunity species, ranging over eleven orders of magnitude even in this relatively small metacommunity of size $J_M = 100,000$. For metacommunities of sizes in the billions or trillions, the curves are much higher on the y -axis, and the longevities range over far greater ranges even than those in figure 8.7. These huge differences in persistence time of common and rare metacommunity species have significant consequences for phylogeny, particularly for the number of daughter lineages that each metacommunity species produces.

We can use the theory for the absorbing process of ecological drift for the i th species developed in chapter 4 to count the number of speciation events that are expected in a given lineage before the lineage goes extinct. To do this, we need to add additional states and transition probabilities to the Markovian model for the absorbing process. Let a new species arise as a single individual with probability ν

per birth, and let N_j be the new species. Thus,

$$\Pr\{N_i - 1, N_j = 1 | N_i\} = \left(\frac{N_i}{J}\right)\nu$$

$$\Pr\{N_i - 1, N_j = 0 | N_i\} = \left(\frac{N_i}{J}\right)(1 - \nu)\left(\frac{J - N_i}{J - 1}\right)$$

$$\Pr\{N_i | N_i\} = \left(\frac{N_i}{J}\right)(1 - \nu)\left(\frac{N_i - 1}{J - 1}\right) \\ + \left(\frac{J - N_i}{J}\right)\left(\frac{J - N_i + 1}{J - 1}\right)$$

$$\Pr\{N_i + 1 | N_i\} = \left(\frac{J - N_i}{J}\right)\left(\frac{N_i}{J - 1}\right)$$

In this system, monodominance is no longer an absorbing state because the monodominant species can be invaded by the new species. In the $D = 1$ case, note that a new species can be generated only when the i th species declines in abundance by one individual (the point mutation version of speciation). The first equation is the transition probability for a speciation event. The second equation is the transition probability for a decline in abundance but no speciation event in the i th species.

In these equations, we are counting only new species whose ancestor is species i . We must do this counting while simultaneously allowing the i th species to take its zero-sum random walk, so we need to keep track of the abundance of the i th species as speciation events occur. This means that we need $2J + 1$ states that record both the current abundance of the i th species as well as whether a speciation event has or has not just occurred. There are only J additional states, not $J + 1$, because the absorbing state of zero abundance cannot produce a new species.

We can now solve for the number of times that the i th species gives rise to a new species as a function of the abundance of species i . Once again ignoring the monodominant state, the number of daughter species Λ produced over the

METACOMMUNITY BIODIVERSITY

evolutionary lifetime of species i having initial abundance N_i is

$$\Lambda(N_i) = \frac{\nu}{J_M} \sum_{k=1}^{N_i-1} k^2 \frac{\binom{J_M}{k}}{\binom{J_M-2}{k-1}} + \frac{N_i \nu}{J_M} \sum_{k=N_i}^{J_M-1} k \frac{\binom{J_M}{k}}{\binom{J_M-2}{k-1}}.$$

From $\Lambda(N_i)$, it is clear that the likelihood that any given metacommunity species will produce a daughter species is small to very small. This is because the speciation rate and metacommunity size enter $\Lambda(N_i)$ as the *ratio* of a very small number to a very large number, ν/J_M . Nevertheless, because of the enormous differences in longevities among species, common species will produce most of the new species that do originate. We have already demonstrated this result and seen the behavior of $\Lambda(N_i)$ in figure 8.3.

These results have immediate application to the reconstruction of phylogenies. The appearance of clades under the unified theory is quite different from those that arise under the models of Raup et al. (1973) and Nee et al. (1994). Consider, for example, a cladogram produced by the Raup et al. model (fig. 8.8). In this model, recall that if the probability of birth λ is greater than the probability of death μ of a lineage, then the usual outcome is exponential growth of the number of lineages. This is true for the example illustrated, a case in which λ was set equal to $9 \cdot 10^{-5}$ and μ was set to $4 \cdot 10^{-5}$. In figure 8.8, I have illustrated Nee et al.'s case 4 (case 2 in fig. 8.1), showing only the "censored" phylogeny of lineages surviving to the present (right edge). Note that case 2 cladograms always make it appear that the speciation rate is increasing with time. This is a visual artifact of pruning out all extinct lineages. In actuality, the probability of lineage birth and death has remained constant through time. Nevertheless, there is a real geometric increase in the number of lineages.

I now illustrate two cladograms produced by the unified theory, under the point mutation mode of speciation. The first is an example of a cladogram produced for a value of



FIG. 8.8. A typical case 2 cladogram produced by the stochastic phylogeny models of Raup et al. (1973) and Nee et al. (1994), which assume constant per lineage probabilities of birth λ and death μ . In this example, $\lambda = 9 \times 10^{-5}$ and $\mu = 4 \times 10^{-5}$. Note the geometric increase in number of lineages. This is partly an artifact of pruning out all extinct lineages.

the fundamental biodiversity number θ of 100, again showing only extant lineages (fig. 8.9). Several differences from figure 8.8 are immediately apparent. First, lineages and biodiversity are much more concentrated at the present time and in the relatively recent past. Once again, this is a case 2 reconstructed phylogeny, so part of this effect is due to pruning out the extinct lineages. But mostly this effect is due to the fact that new lineages arise as rare species that have a very high probability of rapid extinction. Second, there are many more old lineages extending much farther back in time than in the Raup model (figure 8.8). This is not because the clade is polyphyletic, but because a number of very abundant lineages are very long lived. All share a common ancestor and diverged from it long before the time

METACOMMUNITY BIODIVERSITY




FIG. 8.9. A typical case 2 cladogram produced by the unified theory for a fundamental biodiversity number $\theta = 100$. This is for the point mutation mode of speciation. Note the concentration of most diversity at the present and recent past. This is partly due to the pruning out of extinct lineages. However, in the unified theory under point mutation speciation, the diversity really is more concentrated near the present because of rapid extinction of new but rare species.

horizon shown in figure 8.9. Third, some clades are far more bushy and speciose than others. This is because very abundant metacommunity species are ancestors of more daughter lineages than less abundant lineages. This happens for two reasons. First, speciation occurs on a per birth basis, not a per lineage basis, and many more births occur per unit time in common species than in rare ones. Therefore, common species are more often the progenitors of new species. Second, by far the most important quantitative factor is that very common metacommunity species are extremely long lived and old (quite literally “older than the hills”). For this reason they are sources for new species in the long run far more than short-lived rare species. The final difference to

note is that a steady-state biodiversity is achieved in the phylogenetic clade. Even though it looks like diversity is increasing, the present-time biodiversity is actually in equilibrium because the rate of speciation is in stochastic balance with the rate of extinction at all times. Reading backwards in time from the present (right to left in fig. 8.9), after most of the newly arisen species go extinct, there is a much slower decay of biodiversity in the clade. This is because these ancient lineages are the very abundant metacommunity species that are very resistant to extinction.

The biodiversity equilibrium achieved in phylogenetic clades is strongly dependent on the biodiversity number θ as well as on the mode of speciation. I defer consideration of the phylogenetic pattern that arises under the random fission mode until I discuss speciation later in this chapter. However, I now illustrate what happens when we use a smaller value of the fundamental biodiversity number, $\theta = 10$. From the smaller θ , we should expect a reduction in equilibrium clade diversity, and this is in fact what happens (fig. 8.10). There are at least two other interesting consequences. First, because of zero-sum dynamics, we know that the common species in a lower diversity metacommunity will be more abundant than in a species-rich metacommunity. Therefore, *the proportion of long-lived lineages out of all living lineages in the clade should be higher in metacommunities with a smaller θ* . This is because new species arise more slowly, and previously extant species are therefore, on average, more abundant. Second, a corollary conclusion is that *the average age of most lineages will be older in the metacommunity with a smaller θ* . Neither of these predictions emerges from the Raup et al. (1973) models.

I conclude this section by discussing a sampling issue that is similar to the "veil line" problem posed by Preston (1948, 1962) in the sampling of relative species abundance (chapter 2). What happens to the apparent phylogeny if we do not find all the species? This is an important problem

METACOMMUNITY BIODIVERSITY




FIG. 8.10. A typical case 2 cladogram produced by the unified theory, for a fundamental biodiversity number $\theta = 10$. This is for the point mutation mode of speciation. Note that the average age of species increases with a reduction in θ , and the proportion of long-lived lineages increases. These effects occur because the species in a species-poor metacommunity are more abundant than in a species-rich community because of zero-sum dynamics.

not only for estimating modern living biodiversity on Earth, but also for the classical problem in paleontology of estimating the number of species that have ever lived and therefore the true speciation rate. The essential point to make, which was Preston's point as well, is that we are not equally likely to discover each and every extant lineage or species. Common and widespread species will be found, but as species become rarer and rarer, they are more and more likely to be overlooked. Thus, the sampling of species—both in the modern world and in the fossil record—is heavily biased toward finding the spatially and temporally extremely abundant species, a well-known problem to paleontologists and ecologists alike. This bias is a huge problem for the fossil record, but it is also a severe problem for contemporary biodiversity studies.

I illustrate this sampling problem by showing the cladogram that one obtains if the rare species are not known.




FIG. 8.11. Same phylogeny as in figure 8.9, with $\theta = 100$, except that all rare lineages have been pruned out. This increases the apparent age of species and reduces the apparent diversity of clades at the present.

In figure 8.11, I present the same phylogeny that was shown in figure 8.9, except that now all of the rare species (<32 individuals in this numerical example) have been eliminated. Two things are immediately apparent from figure 8.11. First, the terminal bushiness of the clades at the present time is much reduced. This is because most of the rare species are recent in origin. Fewer old species are rare because rare species do not tend to persist. Second, the average age of lineages appears to be older. This age bias is because the species left in the clade after removing the rare species are more resistant to extinction and therefore more persistent. An important conclusion from the unified theory, therefore, is that *a sampling bias toward common species causes an upward bias in the estimate of the mean age of species*. As rarer and rarer species are collected from a given clade, the estimated mean age of species in the clade will inevitably decrease. The neutral theory affords us a direct method to

quantify the magnitude of this bias. We know the theoretical longevities of metacommunity species from the distribution of $\Omega(N_i)$, so we can simply compute the expected mean lifespans for any threshold abundance N_i and above.

The importance of understanding this bias is more than simply estimating the ages of lineages; it also leads us to a systematic underestimation of the speciation rate. I will have more to say on this topic below.

Thus far I have not directly addressed the speciation process and the deeper meaning of biodiversity itself, namely: What is it that we are really enumerating when we count species in the metacommunity or construct a phylogeny? The unified neutral theory gives a surprising and potentially challenging answer to this question that generates a whole array of additional questions for testing in the future. In chapters 2 and 3, I discussed the almost infinite regress of rarity found in observed relative species abundance distributions. As sample sizes increase, the rarest species become ever rarer relative to the commonest species, so that the range of abundances grows larger and larger. We now know that this is a necessary result from the equilibrium theory of metacommunity biodiversity presented in chapter 5. The unavoidable conclusion is that *the more extensively and finely we look, the more biodiversity we find.*

The second conclusion we can draw from observations of biodiversity is that if we lump species into related clades, such as species into genera and genera into families, *the qualitative pattern of abundance and diversity among the higher taxa remains fundamentally unchanged.* We have already illustrated this phenomenon for families of birds (fig. 8.4). A second example showing this pattern is the distribution of plant species per family on three different spatial scales: BCI, Panama, and the entire world (fig. 8.12). These distributions are all very much like the zero-sum multinomials that characterize the distribution of individuals per species (chapters 3, 5),

CHAPTER EIGHT

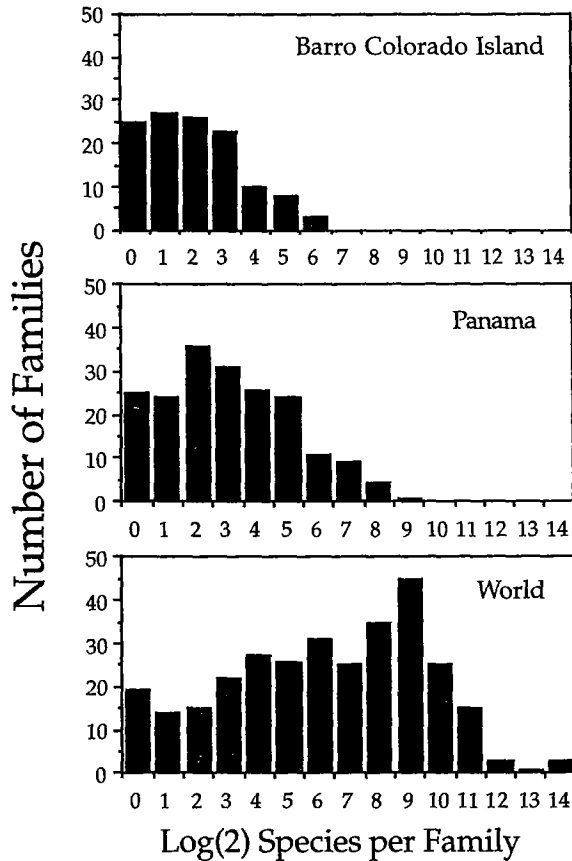


FIG. 8.12. The distribution of species per family of flowering plants, as a Preston-type plot on three different spatial scales—BCI, all of Panama, and the world. Shown is the veil line sampling phenomenon first identified by Preston, and also the long tail of species-poor families, very reminiscent of the zero-sum multinomial distribution of relative species abundances.

for example, the distribution of the abundances of breeding birds of Great Britain (fig. 2.6).

This self-similarity of the organization of biodiversity on different scales of taxonomic resolution leads potentially to a radical departure from the typological view of species because it implies that *biodiversity is intrinsically and fundamentally fractal*. The unified neutral theory provides strong theoretical support for the hypothesis of the fractal nature of biodiversity. Phylogenetic clades generated by the theory have fractal geometry. To see this, perform a sampling of a given clade by counting the number of surviving

METACOMMUNITY BIODIVERSITY




FIG. 8.13. Sampling of a phylogeny produced by the unified neutral theory for measuring its fractal dimension. The number of lineages is counted at exponentially deeper sampling times into the past, and then the logarithm of the number of lineages is plotted against the logarithm of the time depth. This phylogeny was produced by a fundamental biodiversity number θ of 100.

lineages at exponentially increasing sample times into the past (fig. 8.13). Then plot the log of the number of lineages as a function of the log of the time depth of the sampling points on the phylogeny. If these cladograms are self-similar and homogeneous fractals, then these log-log plots should be linear with a slope $-D$, which is the fractal dimension of the phylogeny. The phylogenies predicted by the theory are not only fractal, but the fractal dimension of the phylogeny bears a functional relationship to the fundamental biodiversity number, θ , under the point mutation mode of speciation (fig. 8.14). Under the random fission mode of speciation, the fractal dimension is jointly related both to the speciation rate ν and the size of the metacommunity J_M .




FIG. 8.14. The fractal geometry of phylogenies and of biodiversity under the unified neutral theory. Figure shows the relationship between the log of number of lineages in a clade and the log of the time since the present that the clade is sampled. The fractal dimension of the phylogeny is functionally related to the fundamental biodiversity number θ . Each line is fit to the data from one phylogeny, for $\theta = 100$, 10 or 1. They are not ensemble means of many runs. I chose to illustrate single runs to show the precision of the relationship for single phylogenies. The fractal dimension of a phylogeny decreases with decreasing θ .

This is because these parameters become separated in their effect on metacommunity diversity in the theory for random fission speciation (see below).

If the neutral theory's assertion is correct that biodiversity is fundamentally a fractal, then this conclusion has implications for the way that we think about and count species. If the self-similarity of biodiversity extends smoothly all the way down to the individual level, then the fractal geometry of biodiversity would be homogeneous at all taxonomic scales. But if this is so, where is the biodiversity signature of the existence of species? A signature of species would consist of evidence that the fractal geometry of biodiversity changes and that a different scaling region exists below the species

METACOMMUNITY BIODIVERSITY

level. A change in the scaling region would imply that the processes generating diversity are fundamentally different above and below the species level. Such a change in scaling region would be manifest as a "kink" or bend in the lines describing the fractal geometry of biodiversity (fig. 8.14) at the species level. If no change in slope occurs at the species level, however, then the geometry of biodiversity will provide no evidence for (or against) the special nature of species, even if species have biological reality in other terms, such as in terms of genetic isolation.

The unified neutral theory is silent on the existence of species and their nature. The theory predicts that biodiversity should be a homogeneous fractal with one scaling domain all the way down to the individual level. If biodiversity is a perfectly homogeneous fractal, and if there are no kinks in the fractal geometry at the species level, this implies that a satisfactory answer to the question—how many species are there?—cannot really be answered except operationally. It requires a definition of the scale of aggregation of biodiversity that we call species. This scale is unlikely to be totally arbitrary, even if diversification below the species level is continuous down to the individual level, as the theory asserts. However, this species scale will then be determined less by the true fractal nature of biodiversity, than by the average difference that separates whatever taxonomists identify as "good" species. But whatever we decide is a "good" species, an empirical question remains, namely: To what extent is biodiversity a homogeneous fractal all the way up and down the ladder of taxonomic scale? Even if biodiversity is truly fractal, as the neutral theory asserts, it does not have to have the same self-similarity across all scales. This question and its ramifications are likely to be very fruitful areas for research by ecologists and evolutionary geneticists for years to come.

Up to this point in the book I have explored the neutral theory mainly under the assumption of one particular

CHAPTER EIGHT

mode of speciation, a mode in which new species arise like rare point mutations. I now consider some consequences of changing the mode of speciation to one in which species arise by the random fission of an ancestral species into two daughter species. The decision to explore other modes of speciation was made relatively late in the writing of this book (with the help of my graduate seminar in biogeography), so I have not yet developed the theory for other modes of speciation as fully as would be desirable. A more complete neutral theory for alternative modes of speciation must therefore be set aside for later work. Nevertheless, my results so far suggest that pursuing the implications of different modes of speciation for biodiversity and biogeography would be a valuable and important exercise.

The random fission mode of speciation is an attempt to capture the essence of allopatric speciation (Mayr 1963). The Mayrian concept of species is one of populations reproductively isolated by a complex of pre- and postmating isolating mechanisms. These barriers to the free exchange of genes are hypothesized to arise first in allopatry as correlated responses to selection in the different environments experienced in allopatry. These barriers are believed to be reinforced by selection against hybrids once secondary contact between the populations is established. Partial or total reproductive isolation between species impedes the free flow of genetic information, and this should leave a signature in the fractal geometry of biodiversity.

Whether we fully accept this model of speciation, we can still ask: What will the biodiversity signature of Mayrian species be? Specifically, what impact would the allopatric origination model have upon the metacommunity equilibrium relative species abundance distribution? If we assume that allopatry is required for speciation, then one of the major differences between this mode of speciation and the point mutation mode studied earlier is that incipient allopatric species would usually start their existence at

some high or moderately high initial abundance. The consequences of being common at origination are potentially profound for metacommunity biodiversity because new species will not go extinct as rapidly as they do under the point mutation mode of speciation. Alternatively, it may be that most allopatric speciation events occur in small peripheral isolates, in which case the random fission model may actually be less appropriate than the point mutation model, even though the speciation event occurred in allopatry. From the theoretical point of view, the critical issue is not allopatry per se, but the size of the species population at its origination.

From the perspective of the unified theory, the physical nature of the barrier causing the allopatry is unimportant. The essential point is that an ancestral species population is split into two daughter species. I assume that splitting takes place randomly and cuts the ancestral population into two generally unequal fragments, like the single cut of a knife. I assume that this is a random uniform process, equally likely to cut the population into any two sizes summing to the ancestral population size. This is analogous to a geographic barrier randomly dividing the ancestral population. From the point mutation mode I retain the definition of speciation rates on a per birth basis. Thus, an individual is picked at random from the metacommunity and its species is determined. I then let this species undergo random fission into two daughter species. In order to make the bookkeeping analogous to the point mutation mode, I also assume that the ancestral species persists as the larger population of the two daughter species. This avoids an accounting problem of pseudo-extinction—false extinctions counted when a species evolves into its own descendant.

The random fission mode of speciation has a dramatic effect on metacommunity diversity. In figure 8.15, I compare the equilibria for point mutation versus random fission for a metacommunity of size = 10,000 and a fundamental biodiversity number $\theta = 10$. The equilibrium number of species




FIG. 8.15. Comparison of the equilibrium distribution of relative species abundance in the metacommunity for the point mutation mode of speciation versus the random fission mode, for a metacommunity size $J_M = 10,000$, and $\theta = 10$. Means ± 1 standard deviation.

is increased by nearly a factor of six in this numerical example. The reason for this large increase in species richness under random fission is the longer average lifespans of new species due to their larger population sizes at origination.

Although it may not be apparent from the dominance diversity curve shown in figure 8.15, there is also a qualitative change in the functional shape of the distribution. It turns out that the metacommunity distribution under random fission speciation is not the logseries, but rather it is a zero-sum multinomial with an interior mode (fig. 8.16). The interior mode arises because rare species are not at as high a frequency as under the point mutation mode. In that case, the high frequency of rare species was maintained by the continual input of new rare species, all of which originate at extreme rarity as singletons.

Another important difference between the two modes of speciation is that the fundamental biodiversity number θ

METACOMMUNITY BIODIVERSITY




FIG. 8.16. An example of a zero-sum multinomial that arises as the equilibrium distribution of relative species abundance in the metacommunity under the random fission mode of speciation. The interior mode differentiates this distribution from the one predicted by the point mutation mode, which yields a logseries in the infinite J_M limit.

no longer uniquely determines the metacommunity distribution of relative species abundance under the random fission mode. The difference is that both metacommunity size J_M and the speciation rate ν have to be specified separately because their effects are no longer the same. Although the distribution is uniquely determined by J_M and ν , these two parameters are decoupled in the functions yielding the probability distributions of relative species abundance in the metacommunity. This decoupling can be proven analytically by solving for the equilibrium distribution of abundances for small J_M . I have not yet been able to find the analytical solution for arbitrarily large J_M , however.

The results presented in figure 8.17 show that θ does not uniquely determine the metacommunity distribution under the random fission mode. In all four cases, $\theta = 4$, but the metacommunity sizes and speciation rates are both being varied. The reasons why J_M and ν are decoupled in the random fission mode of speciation are easy to discern. When new species do not always arise as singletons, but can be any




FIG. 8.17. Metacommunity dominance-diversity curves under the random fission mode of speciation, showing how changes in metacommunity size and speciation rate (not shown) affect the distribution of relative species abundance, while holding θ constant.

arbitrary abundance $< J_M$, then it matters separately what the speciation rates and metacommunity sizes are. Increasing metacommunity size while decreasing the speciation rate to hold θ constant results in a net increase in equilibrium diversity. This is because the average size of the initial populations at speciation is larger in larger metacommunities, increasing the mean lifespan of species, more than offsetting the effect of reduced speciation. The result is that more and more species are present at equilibrium between speciation and extinction. The relative abundance of the commonest species is reduced as metapopulation size increases for a fixed value of θ . Note the progressive shallowing of the curves as J_M increases in figure 8.17.

This is in marked contrast to the behavior of the metacommunity distribution under point mutation speciation. In this case, when we increase metacommunity size J_M holding θ constant, the qualitative shape of the curve remains the same

METACOMMUNITY BIODIVERSITY




FIG. 8.18. Metacommunity dominance-diversity curves under the point mutation mode of speciation, showing that concurrent changes in metacommunity size and speciation rate that maintain a constant θ also produce the same distribution of relative species abundance. The differences at high rank are due solely to sampling effects. Compare with figure 8.17.

for the common species and species of intermediate abundance (fig. 8.18). Only the rare end of the distributions differ, and this is simply a sampling phenomenon. Species at a given rank abundance have larger population sizes in larger metacommunities, so the rank abundance at which sampling deviations from the logseries become apparent occurs at higher rank abundances in larger communities.

I have not yet had time to explore in any depth which of these distributions is a better fit to the data from diverse taxonomic groups. I offer just one example from the work of John Lynch (1989) on the genus of leopard frogs. It should be noted that the data used for this example are less than ideal because they consist of species range data as surrogates for relative species abundance, which may be problematic. Nevertheless, although Lynch argues for an allopatric origin for the majority of species in this genus of New World

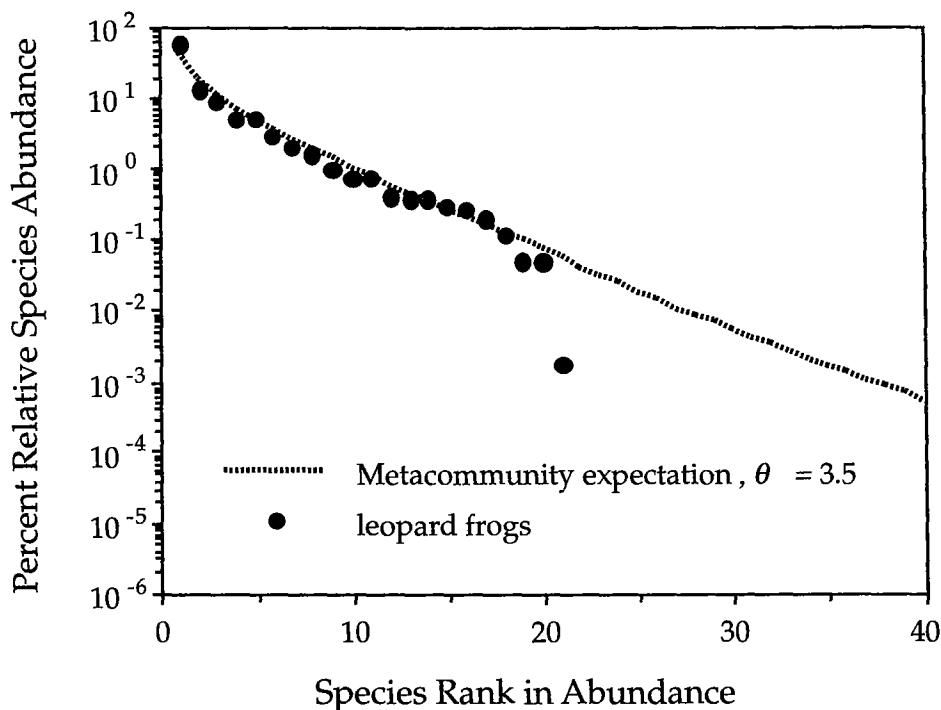


FIG. 8.19. Relationship between percent relative species abundance, as estimated approximately by relative species range sizes, and species rank in abundance for the North American genus of leopard frogs. The estimated value of θ is 3.5. Data from Lynch (1989).

frogs, the data are perhaps more consistent with point mutation speciation (fig. 8.19). If Lynch is correct, then this difference may simply be because it is difficult to distinguish a point mutation mode from speciation events that take place in small peripheral populations, as Lynch proposes for many of the leopard frog species. It is too early to predict what the preponderance of modes of speciation will be once a large sample of different taxa has been analyzed. However, in my limited sampling so far, the point mutation mode has fit more cases than the random fission mode.

The consequence of random fission speciation and a zero-sum multinomial distribution of relative species abundance in the metacommunity for the classical island-mainland problem of MacArthur and Wilson is relatively uneventful. The equilibrium distribution of relative species abundance

on the island or in a local dispersal-limited community remains a zero-sum multinomial distribution. However, there is a further exaggeration of the difference between common and rare species. Rare species on the island are even less frequent than would have been expected from a logseries metacommunity distribution.

These differences raise the exciting possibility that we may be able to use the unified theory to determine what is the dominant mode of speciation in a given metacommunity. Each mode of speciation leaves a different signature in the shape of the metacommunity distribution of relative species abundance. If the metacommunity distribution is best approximated by a logseries, then we can conclude that the point mutation mode is a good approximation to the speciation process. However, if the metacommunity has a zero-sum multinomial distribution with an interior mode and a long tail of ever less frequent rare species (e.g., fig. 8.16), then we can conclude that the random fission mode is a better approximation to the prevailing mode of speciation. For example, one might argue that the distribution of total flowering plant diversity of the world (fig. 8.12) fits the random fission mode. This conclusion is still premature because the graph is based on species per family, which is an imperfect representation of the total abundance of the respective families. As I will document in the next chapter, however, there are also many metacommunity distributions that are completely consistent with the point mutation mode of speciation. The point mutation mode is more consistent with a view of biodiversity as being perfectly self-similar in its fractal geometry at all biological scales of organization.

I do not wish to leave the impression that the mode of speciation must be "either-or." The theory permits a complete continuum between the extreme of point mutation speciation, in which all species are extremely rare at origination, to the random fission mode, in which most species are abundant at origination. If point mutation and random fis-

sion are taken as representing the likely ends of a continuum, then there is a nice conjecture that comes from the neutral theory. We proved in chapter 5 that there was a complete shape continuum between the logseries and the zero-sum multinomial distribution of relative species abundance in local communities, the shape of the distribution depending upon the immigration rate. I cannot prove the following conjecture, but I predict it to be true: *There will also be a complete continuum between the logseries and the zero-sum multinomial in the statistical distributions that fit metacommunity relative abundance distributions. The shape of the fitted distribution will depend upon the relative proportion of point mutation versus random fission speciation events occurring in the particular metacommunity or phylogeny.* In putting forth this conjecture, I am not yet clear in my own mind what the analog to the immigration rate parameter m is (which generates the continuum in the case of local communities), but there must be one. This unknown but anticipated parameter will control the distribution of population sizes of species at origination.

Before leaving the subject of speciation, I return to a question initially posed in chapter 5, namely: What are reasonable rates of speciation, and are these rates consistent with the fitted values of the fundamental biodiversity number, θ ? The number θ typically varies from about 0.1 to 200 or more when fitted to relative abundance data of natural communities, but commonly it is in the rate of 2 to 5 for many taxonomic groups. According to the neutral theory, under the point mutation speciation mode, for the speciation rate to generate these numbers, the speciation rate must be on the order of the inverse of twice the metacommunity size. If the speciation rate is very small, then the metacommunity size will have to be very large. Some of my colleagues have argued that the metacommunity sizes required to get a sufficiently small speciation rate are unrealistic. Or, if one sets the metacommunity sizes small enough to be realistic, then the required speciation rate becomes unrealistically high.

METACOMMUNITY BIODIVERSITY

I don't think we know enough to reach this conclusion. How high is an unrealistic speciation rate? Recall two points from the previous section. First, the fractal nature of biodiversity appears to extend relatively homogeneously all the way down to the individual level. Suppose it is true that the events that ultimately lead to new species begin as small variations among individuals. Second, recall the serious sampling problem associated with finding all rare "species." This sampling issue is a problem in the extreme if the events ultimately leading to speciation occur, and must be detected, at the individual level. Now consider the limiting logseries distribution for the metacommunity under point mutation speciation. This mode is actually a model of infinite diversity in an infinitely large world, analogous to the infinite allele case in population genetics (Ewens 1972).

Of course, in a finite world, there will be a finite number of species (and alleles). Now consider a small sample of this world. There will be a distribution of relative species abundance in this sample. However, to fit the metacommunity logseries distribution, one assumes that the fundamental biodiversity number θ , or Fisher's α , as the case may be, would fit the data for an ever larger sample of the metacommunity, which would show an ever larger number of rare species. In a thought experiment, this could continue until every last individual were censused and classified as to whether it was or was not a new species. It might be difficult in principle to tell whether a particular individual would establish a new species, but suppose we could do so, at least retrospectively (the analog of "Lucy"). Of course, the line of descendants of most individuals is usually very short in an evolutionary sense, even those that might be considered potential founders of new species. So at this level of small-scale resolution of lineage differentiation, what would be the speciation rate? It would, according to the theory, have to be high enough to produce the observed value of the

CHAPTER EIGHT

fundamental biodiversity number θ in finite samples of the metacommunity.

Now consider the sampling problem again. Most of the individual micro-variation leading to new rare species will be totally invisible to the evolutionary biologist and taxonomist. Indeed, the expectation is that by the time a species has become well enough established to be recognized and counted as a "good" species, it is already very abundant and very old. How old is old? This is still an unsolved question. For example, for many years evolutionary biologists thought that the Pleistocene was a major engine of speciation due to constant climate change, but this idea has been largely disproved by recent molecular evidence (Kling and Zink 1997). But even if the average age of most species is several million years, this does not disprove the possibility of a much faster rate operating at the level of individuals, almost all of whose descendant lineages die out rapidly. Following the argument given earlier, our sample of abundant lineages will strongly bias us *against* finding these very short-lived species, and lead to an overestimation of true species ages.

In summary, I think it is premature to dismiss the neutral theory as being inconsistent with measured speciation rates, which are potentially gross underestimates of the rate of origination events happening at the individual level, most of which will go undetected. These underestimates arise because of built-in sampling biases that automatically lead to overestimates of average lineage age.

I conclude this chapter by returning briefly to the subject of species lifespans and the average duration of taxa in the fossil record. In particular, I would like to try to reanalyze the problem of the total biodiversity of Phanerozoic marine genera that ever lived. David Raup (1991) analyzed then unpublished data of Sepkowski on the duration of 17,621 genera of Phanerozoic marine organisms in the fossil record and produced an average survivorship curve for all taxa at

the generic level. Data at the generic level rather than the species level were used to overcome some of the taxonomic problems at the species level for fossil data. Raup fit these data with a proportional hazards model that came out of his neutral theory of cladogenesis, assuming constant probabilities of speciation and extinction per lineage (Raup 1978). Let $S(t)$ be the proportion of genera in the cohort of all genera that survive at least to time t , and let p and q be the rates of origination and extinction, respectively. Then the proportion of genera alive at time t should be given, according to Raup (1991), by

$$S(t) = 1 - \frac{q(e^{(p-q)t} - 1)}{pe^{(p-q)t} - q}.$$

Raup (1991) fit this equation to the observed data, and obtained estimates of genus origination and extinction rates of $p = 0.249$ and $q = 0.250$ per million years, respectively. The data and the fitted $S(t)$ function of time are shown in figure 8.20. Raup (1991) argued that the fact that the fitted speciation and extinction rates are so similar implies that diversity is nearly in equilibrium. However, recall that Raup's model does not, in fact, predict an equilibrium diversity. The only reason Raup was able to fit the proportional hazards function $S(t)$ to the data in figure 8.20 was to assume that diversity is *not*, in fact, in equilibrium. The function $S(t)$ becomes undefined when $p = q$. What is needed is a neutral theory that produces a fit to the data in figure 8.20 when the origination and extinction rates are identical.

The unified theory gives a genuine equilibrium explanation for metacommunity diversity. If we could fit Raup's data to the theory, we could get an estimate of the total number of marine genera that ever lived in the Phanerozoic. Unfortunately, we cannot truly fit the theory to the data because the size and persistence of the metacommunities that supported these marine taxa are not known. Suppose, however, we make the reasonable assumption that the lifespans of the

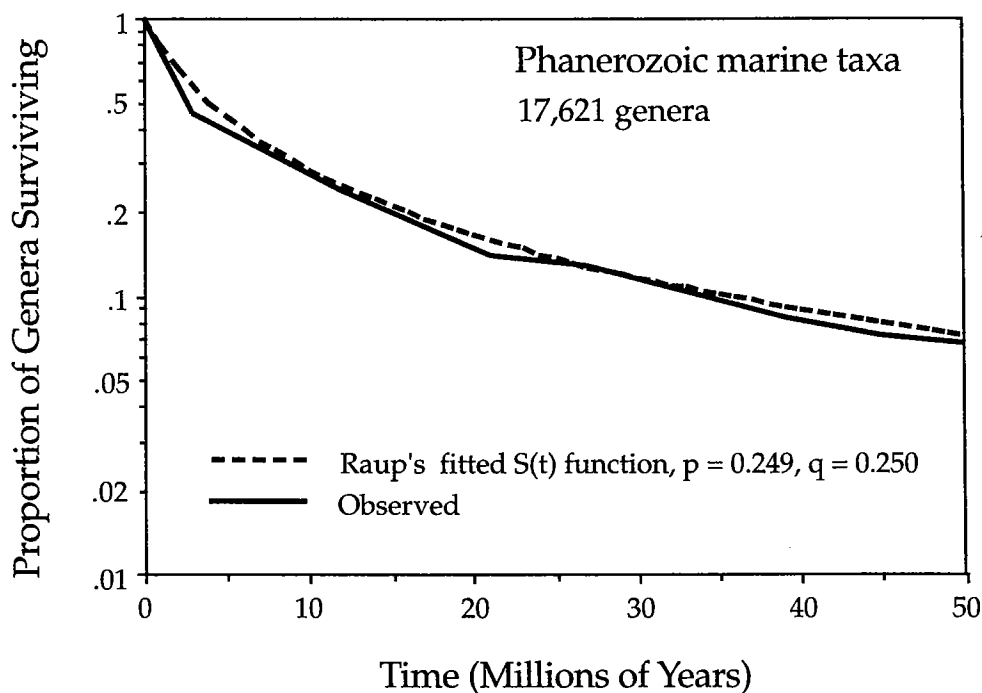


FIG. 8.20. Survivorship curve for 17,621 genera of Phanerozoic marine taxa. The dotted line is the fit of Raup's proportional hazards function, $S(t)$ for a genus origination rate of $p = 0.249$ per million years and an extinction rate $q = 0.250$ per million years. Data were compiled by Sepkowski.

genera are distributed approximately as a zero-sum multinomial. Even this assumption cannot be properly tested because we do not have data on the rarest and shortest-lived taxa. About the best we can do is to fit the data with a lognormal, which will at least give us a conservative estimate of total Phanerozoic marine diversity.

Thus, in the absence of data to fit the zero-sum multinomial distribution, we have to settle for the lognormal, which has almost exactly the same-shaped right-hand tail (for common taxa) as the zero-sum multinomial. Remarkably, the survivorship curve in figure 8.20 is almost perfectly fit by the lognormal distribution. In fitting the lognormal, I used the number of genera going extinct per million years rather than the survivorship curve. Let $S(R)$ be the number of genera going extinct R natural log units of time above or below

the mode. In this case, note that I am not categorizing the number of genera dying into log base two octaves of time. Let S_0 be the modal number of genera dying. Then the equation for the lognormal after taking logs is

$$\ln[S(R)] = \ln(S_0) - a^2 R^2,$$

where a is a fitted constant. Since we have not observed the mode of the distribution in figure 8.20, we also need to estimate a parameter giving the position of the mode on the x -axis. We can fit the lognormal to the data in figure 8.20, and when we find the best fit by maximum likelihood, the fit is excellent (fig. 8.21). There is a slight deviation of the points for short extinction times. However, short extinction times are likely to have greater percentage measurement errors, in any case (Raup 1991), so these deviations, already small, are rendered completely unimportant. From the fitted lognormal equation, we find that the mode of the distribution is located at \log_e (million years) = -1.5 , which is equal to 223,130 years. The number of genera at the mode is $\exp(8.659) = 5,762$. The value of a^2 is 0.187. The parameter a^2 is related to the variance of the lognormal by $a^2 = 1/2\sigma^2$, so the standard deviation of the lognormal in figure 8.21 is $\sigma = 1.634 \log_e$ (million years).

The graph in figure 8.21 only shows the right-hand tail of the lognormal for those genera that lived long enough or were abundant enough in the metacommunity to be fossilized and discovered by paleontologists. All shorter-lived or less-abundant genera that died out in timespans shorter than a million years are not represented in the known fossil record. However, now that we have a lognormal fit to the known fossil record, we can plot the full lognormal curve, below the veil line, which in figure 8.20 is set at one million years. Figure 8.22 shows the full lognormal with its modal position at 223,130 years. The longest-lived taxa have an approximate age of $10^6 e^4$ years, or 54.6 million years. The




FIG. 8.21. Fit of Raup's survivorship curve to the lognormal. Only the right-hand tail of the distribution is "visible" from the fossil record, so we needed to estimate the position of the mode, which is "veiled," *sensu* Preston (1948, 1962). Maximum likelihood estimates placed the mode at 223,130 years. the modal number of genera dying is $\exp(8.659) = 5762$. And the standard deviation of the lognormal is 1.634 ln (million years).

shortest-lived taxa are estimated to live $10^6 e^{-6}$ years, or only 2479 years.

With the full lognormal in hand, we can now estimate the total number of Phanerozoic genera that have ever lived, as well as the percentage of genera that have been found and that have never been found. From the position of the mode and the veil line, which lies 0.918 standard deviation units to the right of the mode, we can estimate the percentage of genera that lived but have no known fossil record: 82%. The remaining 18% of all genera constitute the known fossil record of Phanerozoic marine taxa.

Thus, the unified theory's estimate of the total number of genera from the Phanerozoic era that ever lived is approximately 100,000 (actual estimate is 97,894 genera). If these

METACOMMUNITY BIODIVERSITY




FIG. 8.22. Lognormal for the distribution of number of genera going extinct per \log_e million years, fitted to the survival data of Raup (1991), from data compiled by Sepkowski. The relative position of the mode and the veil line give an estimate of 97,894 genera that ever lived in the Phanerozoic, of which 82% are estimated not yet to have been found.

numbers are approximately correct, then the fraction of all taxa that have been found is considerably higher than estimates of between 0.1% and 1%. If, however, the true distribution in figure 8.21 is an asymmetrical zero-sum multinomial, then we may have considerably underestimated the area under the left-hand tail. It is impossible to say what the shape of the left-hand tail might be until we have better relative abundance data for the known fossil taxa. There we must leave it.

SUMMARY

1. Existing neutral models of phylogeny and phylogenetic reconstruction all assume that lineages have the same probabilities of going extinct and of speciating.

CHAPTER EIGHT

2. This assumption flies in the face of strong evidence, both from ecological as well as paleobiological studies, showing that widespread and abundant lineages on average are much more resistant to extinction and have much longer lifespans.
3. Factoring in lineage abundance gives more realistic patterns of phylogeny that are qualitatively more consistent with observed patterns. Among other things, the theory predicts that most rare and endemic species are relatively recent in origin, whereas widespread abundant species are expected to be older and closer to stem ancestors.
4. The neutral theory predicts that phylogenetic clades are fractal and self-similar on all taxonomic scales from individuals on up to higher taxa. The fractal dimension of the clade is functionally related to the fundamental biodiversity number.
5. The fractal nature of clades suggests that biodiversity itself is fractal. This raises questions about the meaning of species and how to measure the speciation rate.
6. The theory predicts that the mode of speciation will leave a signature in the distribution of relative species abundance. Under point mutation speciation, the metacommunity relative abundance distribution is asymptotically logseries, whereas under random fission speciation, the distribution is zero-sum multinomial.
7. With some assumptions, the neutral theory can be used to estimate the total Phanerozoic marine diversity that ever lived; the estimate is that only about 18% of the nearly 100,000 marine genera that ever lived have been found so far.